arXiv:1603.06216v1 [cs.SY] 20 Mar 2016

Skew-t Inference with Improved Covariance Matrix Approximation

Henri Nurminen, Tohid Ardeshiri, Robert Piché, and Fredrik Gustafsson,

***Abstract*—Filtering and smoothing algorithms for linear discrete-time state-space models with skew-t distributed measurement noise are presented. The proposed algorithms improve upon our earlier proposed filter and smoother using the mean field variational Bayes approximation of the posterior distribution to a skew-t likelihood and normal prior. Our simulations show that the proposed variational Bayes approximation gives a more accurate approximation of the posterior covariance matrix than our earlier proposed method. Furthermore, the novel filter and smoother outperform our earlier proposed methods and conventional low complexity alternatives in accuracy and speed.**

***Index Terms*—skew t, skewness, t-distribution, robust filtering, Kalman filter, variational Bayes, RTS smoother, truncated normal distribution**

I. Introduction

Asymmetric and heavy-tailed noise processes are present in many inference problems. In radio signal based distance estimation [1]–[3], for example, obstacles cause large positive errors that dominate over symmetrically distributed errors from other sources [4]. The skew t-distribution [5]–[7] is the generalization of the t-distribution that has the modeling flexibility to capture both skewness and heavy-tailedness of such noise processes. To exemplify this, Fig. 1 illustrates the contours of the likelihood function for three independent range measurements where some of the measurements are positive outliers. In this example, skew-t, t, and normal likelihoods are compared. The skew-t likelihood gives a more realistic spread of the probability mass than the normal and t likelihoods.

Filtering and smoothing algorithms for linear discrete-time state-space models with skew-t measurement noise using a variational Bayes (VB) method are presented in [8]. This filter is applied to indoor localization with real ultra-wideband data in [9]. This letter proposes improvements to the filter and smoother proposed in [8]. Analogous to [8], the measurement noise is modeled by the skew t-distribution, and the proposed filter and smoother use a VB approximation of the posterior. However, the main contributions of this letter are (1) a new factorization of the approximate posterior distribution, (2) the application of an existing method for approximating the statistics of a truncated multivariate normal distribution (TMND), and (3) a proof of optimality for a truncation ordering in approximation of the moments of the TMND. A TMND is a multivariate normal distribution whose support is restricted (truncated) by linear constraints and that is re-normalized to integrate to unity. The aforementioned contributions improve the estimation performance by reducing the covariance underestimation common to most VB inference algorithms [10, Chapter 10]. To our knowledge, VB approximations have been applied to the skew t-distribution only in our work [8], [9] and by Wand et al. [11].

H. Nurminen and R. Piché are with the Department of Automation Science and Engineering, Tampere University of Technology (TUT), PO Box 692, 33101 Tampere, Finland (e-mails: henri.nurminen@tut.fi, robert.piche@tut.fi). H. Nurminen receives funding from TUT Graduate School, the Foundation of Nokia Corporation, and Tekniikan edistämissäätiö.

T. Ardeshiri and F. Gustafsson are with the Department of Electrical Engineering, Linköping University, 58183 Linköping, Sweden, (e-mails: tohid@isy.liu.se, fredrik@isy.liu.se). T. Ardeshiri receives funding from Swedish research council (VR), project scalable Kalman filters.

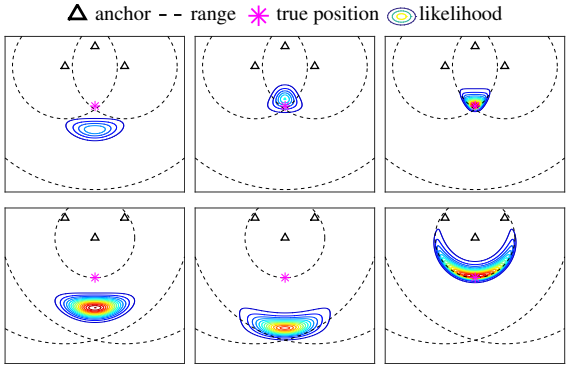

Fig. 1. The contours of the likelihood function for three range measurements for the normal (left), t (middle) and skew-t (right) measurement noise models are presented. The t and skew-t likelihoods can handle one outlier (upper row), while only the skew-t model can handle the two positive outlier measurements (bottom row) due to its asymmetry. The likelihoods' parameters are selected such that the first two moments of the normal, t and skew-t PDFs coincide.

II. Problem formulation

Consider the linear and Gaussian state evolution model

$$x_{k+1} = Ax_k + w_k, \qquad w_k \overset{\text{iid}}{\sim} \mathcal{N}(0, Q), \tag{1a}$$
$$p(x_1) = \mathcal{N}(x_1; x_{1|0}, P_{1|0}), \tag{1b}$$

where $\mathcal{N}(\cdot; \mu, \Sigma)$ denotes a (multivariate) normal PDF with mean μ and covariance matrix Σ; $A \in \mathbb{R}^{n_x \times n_x}$ is the state transition matrix; $x_k \in \mathbb{R}^{n_x}$ indexed by $1 \leq k \leq K$ is the state to be estimated with initial prior distribution (1b), where the subscript "$a|b$" is read "at time a using measurements up to time b". Further, consider the measurements $y_k \in \mathbb{R}^{n_y}$ to be governed by the measurement equation

$$y_k = Cx_k + e_k, \qquad [e_k]_i \overset{\text{iid}}{\sim} \mathrm{ST}(0, R_{ii}, \Delta_{ii}, \nu_i), \tag{2}$$

where the measurement noise distribution is a product of independent univariate skew t-distributions. This model is justified in applications where one-dimensional data from different sensors can be assumed to have statistically independent noise [9]. The PDF and the first two moments of the skew t-distribution can be found in [9] and [12], respectively.

The model (2) admits the hierarchical representation

$$y_k | x_k, u_k, \Lambda_k \sim \mathcal{N}(Cx_k + \Delta u_k, \Lambda_k^{-1} R), \tag{3a}$$
$$u_k | \Lambda_k \sim \mathcal{N}_+(0, \Lambda_k^{-1}), \tag{3b}$$
$$[\Lambda_k]_{ii} \sim \mathcal{G}(\tfrac{\nu_i}{2}, \tfrac{\nu_i}{2}), \tag{3c}$$

where $R \in \mathbb{R}^{n_y \times n_y}$ is a diagonal matrix of which the square roots of the diagonal elements, $\sqrt{R_{ii}}$, are the spread parameters of the skew t-distribution in (2); $\Delta \in \mathbb{R}^{n_y \times n_y}$ is a diagonal matrix whose diagonal elements Δ_{ii} are the shape parameters; $\nu \in \mathbb{R}^{n_y}$ is a vector whose elements ν_i are the degrees of freedom; $C \in \mathbb{R}^{n_y \times n_x}$ is the measurement matrix; $\{w_k \in \mathbb{R}^{n_x} | 1 \leq k \leq K\}$ and $\{e_k \in \mathbb{R}^{n_y} | 1 \leq k \leq K\}$ are mutually independent noise sequences; the operator $[\cdot]_{ij}$ gives the (i, j) entry of its argument; Λ_k is a diagonal matrix with a priori independent random diagonal elements $[\Lambda_k]_{ii}$. Also, $\mathcal{N}_+(\mu, \Sigma)$ is the TMND with closed positive orthant as support, location parameter μ, and squared-scale matrix Σ. Furthermore, $\mathcal{G}(\alpha, \beta)$ is the gamma distribution with shape parameter α and rate parameter β. Models where the measurement noise components are vector-valued with independently multivariate skew-t distributed noises [5]–[7], [13]–[15] require only a straightforward modification to the update of the approximate posterior of Λ_k in the proposed filtering and smoothing algorithms.

Bayesian smoothing means finding the smoothing posterior $p(x_{1:K}, u_{1:K}, \Lambda_{1:K} | y_{1:K})$. In [8], the smoothing posterior is approximated by a factorized distribution of the form $q_{\,[8]} \triangleq q_x(x_{1:K}) q_u(u_{1:K}) q_\Lambda(\Lambda_{1:K})$. Subsequently, the approximate posterior distributions are computed using the VB approach. The VB approach minimizes the Kullback–Leibler divergence (KLD) $D_{\mathrm{KL}}(q||p) \triangleq \int q(x) \log \frac{q(x)}{p(x)} \mathrm{d}x$ [16] of the true posterior from the factorized approximation. That is, $D_{\mathrm{KL}}(q_{\,[8]} || p(x_{1:K}, u_{1:K}, \Lambda_{1:K} | y_{1:K}))$ is minimized in [8].

The numerical simulations in [8] manifest the covariance underestimation of the VB approach, which is a known weakness of the method [10, Chapter 10]. The aim of this letter is to reduce the covariance underestimation of the filter and smoother proposed in [8] by removing independence approximations of the posterior approximation.

III. Proposed Solution

Using Bayes' theorem, the state evolution model (1), and the likelihood (3), the joint smoothing posterior PDF can be derived as in [8]. This posterior is not analytically tractable. We propose to seek an approximation in the form

$$p(x_{1:K}, u_{1:K}, \Lambda_{1:K} | y_{1:K}) \approx q_{xu}(x_{1:K}, u_{1:K})\, q_\Lambda(\Lambda_{1:K}), \quad (4)$$

where the factors in (11) are specified by

$$\hat{q}_{xu}, \hat{q}_\Lambda = \underset{q_{xu}, q_\Lambda}{\operatorname{argmin}} D_{\mathrm{KL}}(q_{\mathrm{N}} || p(x_{1:K}, u_{1:K}, \Lambda_{1:K} | y_{1:K}))$$

and where $q_{\mathrm{N}} \triangleq q_{xu}(x_{1:K}, u_{1:K}) q_\Lambda(\Lambda_{1:K})$. Hence, $x_{1:K}$ and $u_{1:K}$ are not approximated as independent as in [8] because they can be highly correlated a posteriori [8]. The analytical solutions for $\hat{q}_{xu}$ and $\hat{q}_\Lambda$ are obtained by cyclic iteration of

$$\log q_{xu}(\cdot) \leftarrow \underset{q_\Lambda}{\mathbb{E}} [\log p(y_{1:K}, x_{1:K}, u_{1:K}, \Lambda_{1:K})] + c_{xu} \quad (5a)$$

$$\log q_\Lambda(\cdot) \leftarrow \underset{q_{xu}}{\mathbb{E}} [\log p(y_{1:K}, x_{1:K}, u_{1:K}, \Lambda_{1:K})] + c_\Lambda \quad (5b)$$

where the expected values on the right hand sides are taken with respect to the current q_{xu} and q_Λ [10, Chapter 10] [17], [18]. Also, c_{xu} and c_Λ are constants with respect to the variables $(x_{1:K}, u_{1:K})$ and $\Lambda_{1:K}$, respectively.

Computation of the expectation in (5b) requires the first two moments of a TMND, because the support of $u_{1:K}$ is the non-negative orthant. These moments can be computed using the formulas presented in [19]. They require evaluating the CDF (cumulative distribution function) of general multivariate normal distributions. The MATLAB function `mvncdf` implements the numerical quadrature of [20] in 2 and 3 dimensional cases and the quasi-Monte Carlo method of [21] for the dimensionalities $4-25$. However, these methods can be prohibitively slow. Therefore, we approximate the TMND's moments using the fast recursive algorithm suggested in [22], [23]. The method is initialized with the original normal density whose parameters are then updated by applying one linear constraint at a time. For each constraint, the mean and covariance matrix of the once-truncated normal distribution are computed analytically, and the once-truncated distribution is approximated by a non-truncated normal with the updated moments.

Table I
Optimal Recursive Truncation to the Positive Orthant

Inputs: μ, Σ, and the set of the truncated components' indices $\mathcal{T}$
while $\mathcal{T} \neq \emptyset$ **do**
  $k \leftarrow \operatorname{argmin}_i \{\mu_i / \sqrt{\Sigma_{ii}} \mid i \in \mathcal{T}\}$
  $\xi \leftarrow \mu_k / \sqrt{\Sigma_{kk}}$
  if $\Phi(\xi)$ does not underflow to 0 **then**
    $\epsilon \leftarrow \phi(\xi)/\Phi(\xi)$ ▷ ϕ is the PDF of $\mathcal{N}(0,1)$, Φ its CDF
    $\mu \leftarrow \mu + (\epsilon / \sqrt{\Sigma_{kk}}) \cdot \Sigma_{:,k}$
    $\Sigma \leftarrow \Sigma - ((\xi\epsilon + \epsilon^2)/\Sigma_{kk}) \cdot \Sigma_{:,k} \Sigma_{k,:}$
  else
    $\mu \leftarrow \mu + (-\xi / \sqrt{\Sigma_{kk}}) \cdot \Sigma_{:,k}$ ▷ $\lim_{\xi \to -\infty} (\epsilon + \xi) = 0$ [22]
    $\Sigma \leftarrow \Sigma - (1/\Sigma_{kk}) \cdot \Sigma_{:,k} \Sigma_{k,:}$ ▷ $\lim_{\xi \to -\infty} (\xi\epsilon + \epsilon^2) = 1$ [22]
  end if
  $\mathcal{T} \leftarrow \mathcal{T} \backslash \{k\}$
end while
Outputs: μ **and** Σ; $([\mu, \Sigma] \leftarrow \mathtt{rec_trunc}(\mu, \Sigma, \mathcal{T}))$

The result of the recursive truncation depends on the order in which the constraints are applied. Finding the optimal order of applying the truncations is a combinatorial problem. Hence, we choose a greedy approach, whereby the constraint to be applied is chosen from among the remaining constraints so that the resulting once-truncated normal is closest to the true TMND. By Lemma 1, the optimal constraint in KLD-sense is the one that truncates the most probability. The obtained algorithm with the optimal processing sequence for computing the mean and covariance of a given normal distribution truncated to the positive orthant is given in Table I.

Lemma 1. *Let $p(\mathbf{z})$ be a TMND with the support $\{\mathbf{z} \geq 0\}$ and $q(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mu, \Sigma)$. Then,*

$$\underset{i}{\operatorname{argmin}}\, D_{\mathrm{KL}} \left(p(\mathbf{z}) \,||\, \tfrac{1}{c_i} q(\mathbf{z}) [\![\mathbf{z}_i \geq 0]\!] \right) = \underset{i}{\operatorname{argmin}} \tfrac{\mu_i}{\sqrt{\Sigma_{ii}}}, \quad (6)$$

where μ_i is the ith element of μ, Σ_{ii} is the ith diagonal element of Σ, $[\![\cdot]\!]$ is the Iverson bracket, and $c_i = \int q(\mathbf{z}) [\![\mathbf{z}_i \geq 0]\!] \, \mathrm{d}\mathbf{z}$.

Proof:

$$\begin{aligned} & D_{\mathrm{KL}} \left(p(\mathbf{z}) \,||\, \tfrac{1}{c_i} q(\mathbf{z}) [\![\mathbf{z}_i \geq 0]\!] \right) \\ & \quad \stackrel{+}{=} - \int_0^\infty p(\mathbf{z}) \log(\tfrac{1}{c_i} q(\mathbf{z}) [\![\mathbf{z}_i \geq 0]\!]) \, \mathrm{d}\mathbf{z} \\ & \quad = \log c_i - \int_0^\infty p(\mathbf{z}) \log q(\mathbf{z}) \, \mathrm{d}\mathbf{z} - 1 \stackrel{+}{=} \log c_i, \end{aligned}$$

where $\stackrel{+}{=}$ means equality up to an additive constant. Since c_i is an increasing function of $\frac{\mu_i}{\sqrt{\Sigma_{ii}}}$ the proof follows. ■

The recursion (5) is convergent to a local optimum [10, Chapter 10]. However, there is no proof of convergence available when the moments of the TMND are approximated. In spite of lack of a convergence proof the iterations did not diverge in the numerical simulations presented in section IV.

The derivations for the expectations of (5) are presented in the appendixes. In the smoother, the update (5a) includes

Table II
SMOOTHING FOR SKEW-t MEASUREMENT NOISE

Inputs: A, C, Q, R, Δ, ν, $x_{1|0}$, $P_{1|0}$ and $y_{1:K}$
$A_z \leftarrow \left[\begin{smallmatrix} A & 0 \\ 0 & 0 \end{smallmatrix}\right]$, $C_z \leftarrow [\, C \;\; \Delta \,]$
initialization
$\Lambda_{k|K} \leftarrow I_{n_y}$ for $k = 1 \cdots K$
repeat
  update $q_{xu}(x_{1:K}, u_{1:K})$ *given* $q_\Lambda(\Lambda_{1:K})$
  for k = 1 to K **do**
    $Z_{k|k-1} \leftarrow \mathrm{blockdiag}(P_{k|k-1}, \Lambda_{k|K}^{-1})$
    $K_z \leftarrow Z_{k|k-1} C_z^{\mathrm{T}} (C P_{k|k-1} C^{\mathrm{T}} + \Delta \Lambda_{k|K}^{-1} \Delta^{\mathrm{T}} + \Lambda_{k|K}^{-1} R)^{-1}$
    $\widetilde{z}_{k|k} \leftarrow \left[\begin{smallmatrix} x_{k|k-1} \\ 0 \end{smallmatrix}\right] + K_z (y_k - C x_{k|k-1})$
    $\widetilde{Z}_{k|k} \leftarrow (I - K_z C_z) P_{k|k-1}$
    $[z_{k|k}, Z_{k|k}] \leftarrow \mathtt{rec_trunc}(\widetilde{z}_{k|k}, \widetilde{Z}_{k|k}, \{n_x + 1 \cdots n_x + n_y\})$
    $x_{k|k} \leftarrow [z_{k|k}]_{1:n_x}$, $P_{k|k} \leftarrow [Z_{k|k}]_{1:n_x, 1:n_x}$
    $x_{k+1|k} \leftarrow A x_{k|k}$
    $P_{k+1|k} \leftarrow A P_{k|k} A^{\mathrm{T}} + Q$
  end for
  for $k = K - 1$ down to 1 **do**
    $G_k \leftarrow Z_{k|k} A_z Z_{k+1|k}^{-1}$
    $z_{k|K} \leftarrow z_{k|k} + G_k (z_{k+1|K} - A_z z_{k|k})$
    $Z_{k|K} \leftarrow Z_{k|k} + G_k (Z_{k+1|K} - Z_{k+1|k}) G_k^{\mathrm{T}}$
    $x_{k|K} \leftarrow [z_{k|K}]_{1:n_x}$, $P_{k|K} \leftarrow [Z_{k|K}]_{1:n_x, 1:n_x}$
    $u_{k|K} \leftarrow [z_{k|K}]_{n_x + (1:n_y)}$, $U_{k|K} \leftarrow [Z_{k|K}]_{n_x + (1:n_y), n_x + (1:n_y)}$
  end for
  update $q_\Lambda(\Lambda_{1:K})$ *given* $q_{xu}(x_{1:K}, u_{1:K})$
  for k = 1 to K **do**
    $\Psi_k \leftarrow (y_k - C_z z_{k|K})(y_k - C_z z_{k|K})^{\mathrm{T}} R^{-1} + C_z Z_{k|K} C_z^{\mathrm{T}} R^{-1} + u_{k|K} u_{k|K}^{\mathrm{T}} + U_{k|K}$
    $[\Lambda_{k|K}]_{ii} \leftarrow \frac{\nu_i + 2}{\nu_i + [\Psi_k]_{ii}}$
  end for
until converged
Outputs: $x_{k|K}$ **and** $P_{k|K}$ **for** $k = 1 \cdots K$

Table III
FILTERING FOR SKEW-t MEASUREMENT NOISE

Inputs: A, C, Q, R, Δ, ν, $x_{1|0}$, $P_{1|0}$ and $y_{1:K}$
$C_z \leftarrow [\, C \;\; \Delta \,]$
for k = 1 to K **do**
  initialization
  $\Lambda_{k|k} \leftarrow I_{n_y}$
  repeat
    update $q_{xu}(x_k, u_k) = \mathcal{N}(\left[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}\right]; z_{k|k}, Z_{k|k})$ *given* $q_\Lambda(\Lambda_k)$
    $Z_{k|k-1} \leftarrow \mathrm{blockdiag}(P_{k|k-1}, \Lambda_{k|k}^{-1})$
    $K_z \leftarrow Z_{k|k-1} C_z^{\mathrm{T}} (C P_{k|k-1} C^{\mathrm{T}} + \Delta \Lambda_{k|k}^{-1} \Delta^{\mathrm{T}} + \Lambda_{k|k}^{-1} R)^{-1}$
    $\widetilde{z}_{k|k} \leftarrow \left[\begin{smallmatrix} x_{k|k-1} \\ 0 \end{smallmatrix}\right] + K_z (y_k - C x_{k|k-1})$
    $\widetilde{Z}_{k|k} \leftarrow (I - K_z C_z) P_{k|k-1}$
    $[z_{k|k}, Z_{k|k}] \leftarrow \mathtt{rec_trunc}(\widetilde{z}_{k|k}, \widetilde{Z}_{k|k}, \{n_x + 1 \cdots n_x + n_y\})$
    $x_{k|k} \leftarrow [z_{k|k}]_{1:n_x}$, $P_{k|k} \leftarrow [Z_{k|k}]_{1:n_x, 1:n_x}$
    $u_{k|k} \leftarrow [z_{k|k}]_{n_x + (1:n_y)}$, $U_{k|k} \leftarrow [Z_{k|k}]_{n_x + (1:n_y), n_x + (1:n_y)}$
    update $q_\Lambda(\Lambda_k) = \prod_{i=1}^{n_y} \mathcal{G}\left([\Lambda_k]_{ii}; \frac{\nu_i}{2} + 1, \frac{\nu_i + [\Psi_k]_{ii}}{2}\right)$ *given* $q_{xu}(x_k, u_k)$
    $\Psi_k \leftarrow (y_k - C_z z_{k|k})(y_k - C_z z_{k|k})^{\mathrm{T}} R^{-1} + C_z Z_{k|k} C_z^{\mathrm{T}} R^{-1} + u_{k|k} u_{k|k}^{\mathrm{T}} + U_{k|k}$
    $[\Lambda_{k|k}]_{ii} \leftarrow \frac{\nu_i + 2}{\nu_i + [\Psi_k]_{ii}}$
  until converged
  $x_{k+1|k} \leftarrow A x_{k|k}$
  $P_{k+1|k} \leftarrow A P_{k|k} A^{\mathrm{T}} + Q$
end for
Outputs: $x_{k|k}$ **and** $P_{k|k}$ **for** $k = 1 \cdots K$

a forward filtering step of the Rauch–Tung–Striebel smoother (RTSS) [24] where the first filtering posterior is a TMND. The TMND is approximated as a multivariate normal distribution whose parameters are obtained using the recursive truncation. This approximation enables recursive forward filtering and the use of RTSS's backward smoothing step that gives normal approximations to the marginal smoothing posteriors $q_{xu}(x_k, u_k) \approx \mathcal{N}(\left[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}\right]; z_{k|K}, Z_{k|K})$. After the iterations converge, the variables $u_{1:K}$ are integrated out to get the approximate smoothing posteriors $q_x(x_k) = \mathcal{N}(x_k; x_{k|K}, P_{k|K})$, where the parameters $x_{k|K}$ and $P_{k|K}$ are the output of the skew t smoother (STS) algorithm in Table II. STS can be restricted to an online recursive algorithm to synthesize a filter which is summarized in Table III. In the filter, the output of a filtering step is also a TMND which in analogy to STS is approximated by a multivariate normal distribution to have a recursive algorithm. Using recursive truncation, the TMND is approximated by a normal distribution $q_x(x_k) = \mathcal{N}(x_k; x_{k|k}, P_{k|k})$ whose parameters are the outputs of the skew t filter (STF) algorithm in Table III.

IV. SIMULATIONS

Our numerical simulations use satellite navigation pseudorange measurements of the model

$$[y_k]_i \mid x_k \sim \mathrm{ST}(\|s_i - [x_k]_{1:3}\| + [x_k]_4, 1\,\mathrm{m}, \delta\,\mathrm{m}, 4) \qquad (7)$$

where s_i is the ith satellite's position, $[x_k]_4$ is bias with prior $\mathcal{N}(0, (0.75\,\mathrm{m})^2)$, and δ is skewness parameter. The linearization error is negligible because the satellites are far. The state model is a random walk with process covariance $Q = \mathrm{diag}((q\,\mathrm{m})^2, (q\,\mathrm{m})^2, (0.2\,\mathrm{m})^2, 0)$, where q is a parameter. A satellite constellation of Global Positioning System provided by the International GNSS service [25] is used with 8 measured satellites. The RMSE is computed for $[x_k]_{1:3}$.

A. Computation of TMND statistics

In this subsection we study the computation of the moments of the untruncated components of a TMND. One state and one measurement vector per Monte Carlo replication are generated from the model (7) with $\nu = \infty$ degrees of freedom (corresponding to skew-normal likelihood), prior $x \sim \mathcal{N}(0, \mathrm{diag}(\rho\,\mathrm{m}^2, \rho\,\mathrm{m}^2, (0.22\,\mathrm{m})^2, (0.1\,\mathrm{m})^2))$, and 10 000 replications. The compared methods are recursive truncations with the optimal truncation order (RTopt) and with random order (RTrand), the variational Bayes (VB), and the analytical formulas of [19] using MATLAB function `mvncdf` (MVNCDF). In RTrand any of the non-optimal constraints is chosen at each truncation. VB is an update of the skew t variational Bayes filter (STVBF) [8] where $\overline{\Lambda_k} = I$ and the VB iteration is terminated when the position estimate changes less than 0.005 m or at the 1000th iteration.

Fig. 2 shows distributions of the distance from the estimate of the bootstrap particle filter (PF) with 100 000 samples. The box levels are 5 %, 25 %, 50 %, 75 %, and 95 % quantiles and the asterisks show minimum and maximum values. With small ρ and δ the differences between RTrand, RTopt, and MVNCDF are small. With large δ there are statistically significant differences as the p-values of two-sided Wilcoxon signed rank test in Fig. 2 show. RTopt outperforms RTrand in the cases with high skewness, which reflects the result of Lemma 1. MVNCDF is more accurate than RTopt in the cases with high skewness, but MVNCDF's computational load is roughly 40 000 times that of the RTopt. This justifies the use of recursive truncation approximation.

The approximation of the posterior covariance matrix is tested by studying the normalized estimation error squared (NEES) values [26, Ch. 5.4.2] shown by Fig. 3. If the covariance matrix is correct, the expected value of NEES is

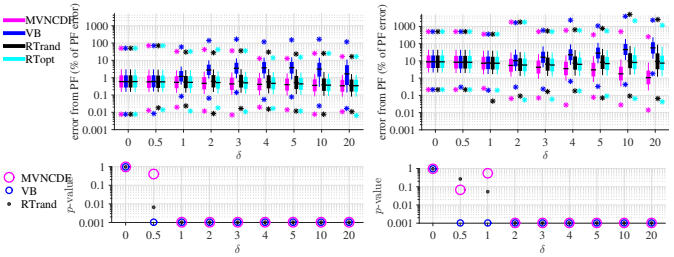

Fig. 2. With large δ values RTopt is closer to PF than RTrand but less accurate than computationally heavy MVNCDF (upper row). p-values of two-sided Wilcoxon signed rank test (bottom row) show that the differences from RTopt are significant with large δ. (left) $\rho = 1^2$, (right) $\rho = 20^2$.

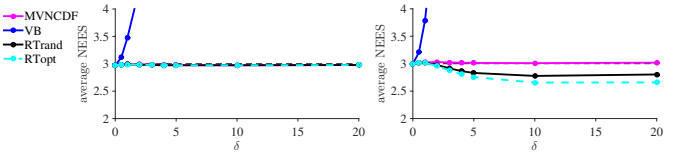

Fig. 3. RTopt's NEES is closest to the optimal value 3, so recursive truncation gives the most realistic covariance matrix. (left) $\rho = 1^2$, (right) $\rho = 20^2$.

the state dimensionality 3 [26, Ch. 5.4.2]. VB gets large NEES values when δ is large, which indicates that VB underestimates the covariance matrix. RTopt and RTrand give NEES values closest to 3, so the recursive truncation provides the most accurate covariance matrix approximation.

B. Skew-t inference

In this section, the proposed skew t filter (STF) is compared with state-of-the-art filters using numerical simulations of a 100-step trajectory. The compared methods are a bootstrap-type PF, STVBF [8], t variational Bayes filter (TVBF) [27], and Kalman filter (KF) with measurement validation gating [26, Ch. 5.7.2] that discards the measurement components whose normalized innovation squared is larger than the χ^2_1-distribution's 99 % quantile. TVBF and KF's parameters are numerically optimized maximum expected likelihood parameters. The results are based on 1000 Monte Carlo replications.

Fig. 4 illustrates the filter iterations' convergence. The figure shows that the proposed STF converges within 5 VB iterations and outperforms the other filters except for PF already with 2 VB iterations. Furthermore, Fig. 4 shows that STF's converged state is close to the PF's converged state in RMSE, and PF can require as many as 10 000 particles to outperform STF. STF also converges faster than STVBF when the process variance parameter q is large. With a small q, STVBF with a small number of VB iterations can give a lower RMSE than the converged STVBF. The reason for this is probably that in the first iterations STVBF accommodates outliers by decreasing the Λ_k estimates, which also affects the covariance, while in the later iterations u_k estimates are increased, which makes the mean more accurate but underestimates the covariance.

Fig. 5 shows the distributions of the RMSE differences from the STF's RMSE as percentages of the STF's RMSE. STF clearly has the smallest RMSE when $\delta \geq 3$. Unlike STVBF, the new STF improves accuracy even with small q, which can be explained by the improved covariance approximation.

Fig. 6 shows the results of a test where the measurement noise in (7) is generated from the histogram distribution of the UWB time-of-flight data set used in [9]. The filters use the maximum likelihood parameters fitted to the data set numerically with the degrees-of-freedom parameters fixed to

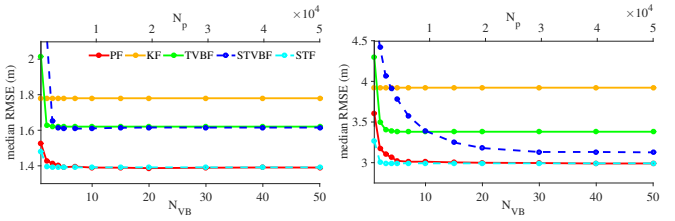

Fig. 4. STF converges in five iterations. The required number of PF particles can be 10.000. (left) $q = 0.5$, $\delta = 5$, (right) $q = 5$, $\delta = 5$.

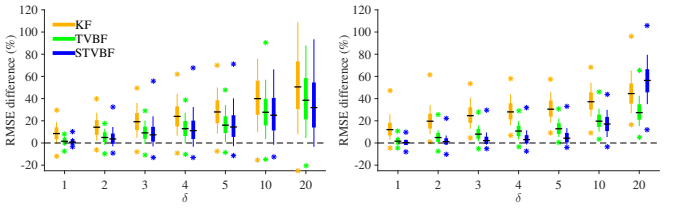

Fig. 5. STF outperforms the comparison methods with skew-t-distributed noise. RMSE differences per cent of the STF's RMSE. The relative differences increase as δ is increased. (left) $q = 0.5$, (right) $q = 5$.

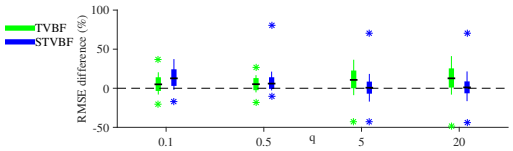

Fig. 6. STF outperforms TVBF and STVBF with UWB noise. RMSE differences per cent of the STF's RMSE.

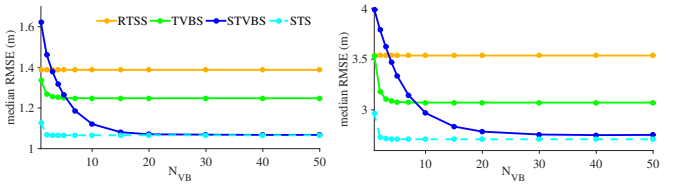

Fig. 7. Five STS iterations give the converged state's RMSE. (left) $q = 0.5$, $\delta = 5$, (right) $q = 5$, $\delta = 5$.

4. The proposed method STF has the lowest RMSE also in this test, which shows that the method is robust to deviations from the assumed distribution and thus usable with real data.

The proposed smoother is also tested with measurements generated from (7). The compared smoothers are the proposed skew t smoother (STS), skew t variational Bayes Smoother (STVBS) [8], t variational Bayes smoother (TVBS) [27], and the RTSS with 99 % measurement validation gating [24]. Fig. 7 shows that STS has lower RMSE than the smoothers based on symmetric distributions. Furthermore, STF's VB iteration converges in five iterations, so it is faster than STVBF.

V. Conclusions

We have proposed a novel approximate filter and smoother for linear state-space models with heavy-tailed and skewed measurement noise distribution. The algorithms are based on the variational Bayes approximation, where some posterior independence approximations are removed from the earlier versions of the algorithms to avoid significant underestimation of the posterior covariance matrix. Removal of independence approximations is enabled by the recursive truncation algorithm for approximating the mean and covariance matrix of truncated multivariate normal distribution. An optimal processing sequence is given for the recursive truncation.

APPENDIX A
DERIVATIONS FOR THE SMOOTHER

We derive the expectations for the iterations of the variational Bayes smoother approximating the joint smoothing density

$$p(x_{1:K}, u_{1:K}, \Lambda_{1:K}|y_{1:K}) \propto p(x_{1:K}, u_{1:K}, \Lambda_{1:K}, y_{1:K}) \tag{8}$$

$$= p(x_1) \prod_{l=1}^{K-1} p(x_{l+1}|x_l) \prod_{k=1}^{K} p(y_k|x_k, u_k, \Lambda_k)\, p(u_k|\Lambda_k)\, p(\Lambda_k) \tag{9}$$

$$= \mathcal{N}(x_1; x_{1|0}, P_{1|0}) \prod_{l=1}^{K-1} \mathcal{N}(x_{l+1}; Ax_l, Q) \cdot \prod_{k=1}^{K} \left\{ \mathcal{N}(y_k; Cx_k + \Delta u_k, \Lambda_k^{-1} R)\, \mathcal{N}_+(u_k; 0, \Lambda_k^{-1}) \prod_{i=1}^{n_y} \mathcal{G}\left([\Lambda_k]_{ii}; \frac{\nu_i}{2}, \frac{\nu_i}{2}\right) \right\} \tag{10}$$

which is approximated by a factorized probability density function (PDF) in the form

$$p(x_{1:K}, u_{1:K}, \Lambda_{1:K}|y_{1:K}) \approx q_{xu}(x_{1:K}, u_{1:K})\, q_\Lambda(\Lambda_{1:K}). \tag{11}$$

The VB solutions $\hat{q}_{xu}$ and $\hat{q}_\Lambda$ can be obtained by cyclic iteration of

$$\log q_{xu}(x_{1:K}, u_{1:K}) \leftarrow \mathop{\mathbb{E}}_{q_\Lambda}[\log p(y_{1:K}, x_{1:K}, u_{1:K}, \Lambda_{1:K})] + c_{xu} \tag{12a}$$

$$\log q_\Lambda(\Lambda_{1:K}) \leftarrow \mathop{\mathbb{E}}_{q_{xu}}[\log p(y_{1:K}, x_{1:K}, u_{1:K}, \Lambda_{1:K})] + c_\Lambda \tag{12b}$$

where the expected values are taken with respect to the current q_{xu} and q_Λ, and c_{xu} and c_Λ are constants with respect to the variables $\left[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}\right]$ and Λ_k, respectively [10, Chapter 10] [17]. This appendix gives the derivations for one iteration of (12). For brevity all constant values are denoted by c. The logarithm of the joint smoothing distribution is

$$\begin{aligned} \log p(x_{1:K}, u_{1:K}, \Lambda_{1:K}, y_{1:K}) &= \log \mathcal{N}(x_1; x_{1|0}, P_{1|0}) + \sum_{l=1}^{K} \log \mathcal{N}(x_{l+1}; Ax_l, Q) \\ &+ \sum_{k=1}^{K-1} \left\{ \log \mathcal{N}(y_k; Cx_k + \Delta u_k, \Lambda_k^{-1} R) + \log \mathcal{N}_+(u_k; 0, \Lambda_k^{-1}) \right\} + \sum_{k=1}^{K} \sum_{i=1}^{n_y} \log \mathcal{G}([\Lambda_k]_{ii}; \tfrac{\nu_i}{2}, \tfrac{\nu_i}{2}), \end{aligned} \tag{13}$$

A. Derivations for q_{xu}

Eq. (12a) gives

$$\begin{aligned} \log q_{xu}(x_{1:K}, u_{1:K}) &= \log \mathcal{N}(x_1; x_{1|0}, P_{1|0}) + \sum_{l=1}^{K-1} \log \mathcal{N}(x_{l+1}; Ax_l, Q) \\ &+ \sum_{k=1}^{K} \mathop{\mathbb{E}}_{q_\Lambda}[\log \mathcal{N}(y_k; Cx_k + \Delta u_k, \Lambda_k^{-1} R) + \log \mathcal{N}_+(u_k; 0, \Lambda_k^{-1})] + c \end{aligned} \tag{14}$$

$$\begin{aligned} &= \log \mathcal{N}(x_1; x_{1|0}, P_{1|0}) + \sum_{l=1}^{K-1} \log \mathcal{N}(x_{l+1}; Ax_l, Q) \\ &- \frac{1}{2} \sum_{k=1}^{K} \mathop{\mathbb{E}}_{q_\Lambda}[(y_k - Cx_k - \Delta u_k)^{\mathrm{T}} R^{-1} \Lambda_k (y_k - Cx_k - \Delta u_k) + u_k^{\mathrm{T}} \Lambda_k u_k] + c \end{aligned} \tag{15}$$

$$\begin{aligned} &= \log \mathcal{N}(x_1; x_{1|0}, P_{1|0}) + \sum_{l=1}^{K-1} \log \mathcal{N}(x_{l+1}; Ax_l, Q) \\ &- \frac{1}{2} \sum_{k=1}^{K} \left\{ (y_k - Cx_k - \Delta u_k)^{\mathrm{T}} R^{-1} \Lambda_{k|K} (y_k - Cx_k - \Delta u_k) + u_k^{\mathrm{T}} \Lambda_{k|K} u_k \right\} + c \end{aligned} \tag{16}$$

$$\begin{aligned} &= \log \mathcal{N}(x_1; x_{1|0}, P_{1|0}) + \sum_{l=1}^{K-1} \log \mathcal{N}(x_{l+1}; Ax_l, Q) \\ &+ \sum_{k=1}^{K} \left\{ \log \mathcal{N}(y_k; Ax_k + \Delta u_k, \Lambda_{k|K}^{-1} R) + \log \mathcal{N}(u_k; 0, \Lambda_{k|K}^{-1}) \right\} + c \end{aligned} \tag{17}$$

$$\begin{aligned} &= \log \mathcal{N}\left(\begin{bmatrix} x_1 \\ u_1 \end{bmatrix}; \begin{bmatrix} x_{1|0} \\ 0 \end{bmatrix}, \begin{bmatrix} P_{1|0} & \mathrm{O} \\ \mathrm{O} & \Lambda_{1|K}^{-1} \end{bmatrix} \right) + \sum_{l=1}^{K-1} \log \mathcal{N}\left(\begin{bmatrix} x_{l+1} \\ u_{l+1} \end{bmatrix}; \begin{bmatrix} A & \mathrm{O} \\ \mathrm{O} & \mathrm{O} \end{bmatrix} \begin{bmatrix} x_l \\ u_l \end{bmatrix}, \begin{bmatrix} Q & \mathrm{O} \\ \mathrm{O} & \Lambda_{l+1|K}^{-1} \end{bmatrix} \right) \\ &+ \log \mathcal{N}\left(y_k; [C \quad \Delta] \begin{bmatrix} x_k \\ u_k \end{bmatrix}, \Lambda_{k|K}^{-1} R \right) + c, \;\; u_{1:K} \geq 0, \end{aligned} \tag{18}$$

where $\Lambda_{k|K} \triangleq \mathbb{E}_{q_\Lambda}[\Lambda_k]$ is derived in Section A-B, and $u_{1:K} \geq 0$ means that all the components of all u_k are required to be nonnegative for each $k = 1 \cdots K$. Up to the truncation of the u components, $q_{xu}(x_{1:K}, u_{1:K})$ has thus the same form as the joint smoothing posterior of a linear state-space model with the state transition matrix $\widetilde{A} \triangleq \left[\begin{smallmatrix} A & \mathrm{O} \\ \mathrm{O} & \mathrm{O} \end{smallmatrix}\right]$, process noise covariance matrix $\widetilde{Q_k} \triangleq \left[\begin{smallmatrix} Q & \mathrm{O} \\ \mathrm{O} & \Lambda^{-1}_{k+1|K} \end{smallmatrix}\right]$, measurement model matrix $\widetilde{C} \triangleq [\, C \; \Delta \,]$, and measurement noise covariance matrix $\widetilde{R} \triangleq \Lambda^{-1}_{k|K} R$. Let us denote the PDFs related to this state-space model with $\widetilde{p}$.

It would be possible to compute the truncated multivariate normal posterior of the joint smoothing distribution $\widetilde{p}\left(\left[\begin{smallmatrix} x_{1:K} \\ u_{1:K} \end{smallmatrix}\right] | y_{1:K}\right)$, and account for the truncation of $u_{1:K}$ to the positive orthant using the recursive truncation. However, this would be impractical with large K due to the large dimensionality $K \times (n_x + n_y)$. A feasible solution is to approximate each filtering distribution in the Rauch–Tung–Striebel smoother's (RTSS [24]) forward filtering step with a multivariate normal distribution by

$$\widetilde{p}(x_k, u_k | y_{1:k}) = \frac{1}{C} \mathcal{N}\left(\begin{bmatrix} x_k \\ u_k \end{bmatrix}; z'_{k|k}, Z'_{k|k} \right) \cdot [u_k \geq 0] \tag{19}$$

$$\approx \mathcal{N}\left(\begin{bmatrix} x_k \\ u_k \end{bmatrix}; z_{k|k}, Z_{k|k} \right) \tag{20}$$

for each $k = 1 \cdots K$, where $[u_k \geq 0]$ is the Iverson bracket notation

$$[u_k \geq 0] = \begin{cases} 1, & \text{if all components of } u_k \text{ are non-negative} \\ 0, & \text{otherwise} \end{cases},$$

C is the normalization factor, and $z_{k|k} \triangleq \mathbb{E}_{\widetilde{p}}\left[\left[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}\right] | y_{1:k}\right]$ and $Z_{k|k} \triangleq \mathrm{Var}_{\widetilde{p}}\left[\left[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}\right] | y_{1:k}\right]$ are approximated using the recursive truncation. Given the multivariate normal approximations of the filtering posteriors $\widetilde{p}(x_k, u_k | y_{1:k})$, by Lemma 2 the backward recursion of the RTSS gives multivariate normal approximations of the smoothing posteriors $\widetilde{p}(x_k, u_k | y_{1:K})$. The quantities required in the derivations of Section A-B are the expectations of the smoother posteriors $x_{k|K} \triangleq \mathbb{E}_{q_{xu}}[x_k]$, $u_{k|K} \triangleq \mathbb{E}_{q_{xu}}[u_k]$, and the covariance matrices $Z_{k|K} \triangleq \mathrm{Var}_{q_{xu}}\left[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}\right]$ and $U_{k|K} \triangleq \mathrm{Var}_{q_{xu}}[u_k]$.

Lemma 2. *Let $\{z_k\}_{k=1}^K$ be a linear–Gaussian process, and $\{y_k\}_{k=1}^K$ a measurement process such that*

$$z_1 \sim \mathcal{N}(z_{1|0}, Z_{1|0}) \tag{21}$$

$$z_k | z_{k-1} \sim \mathcal{N}(A z_{k-1}, Q) \tag{22}$$

$$y_k | z_k \sim \textit{(a known distribution)}. \tag{23}$$

with the standard Markovianity assumptions. Then, if the filtering posterior $p(z_k | y_{1:k})$ is a multivariate normal distribution for each k, then for each $k < K$

$$z_k | y_{1:K} \sim \mathcal{N}(z_{k|K}, Z_{k|K}), \tag{24}$$

where

$$z_{k|K} = z_{k|k} + G_k(z_{k+1|K} - A z_{k|k}), \tag{25}$$

$$Z_{k|K} = Z_{k|k} + G_k(Z_{k+1|K} - A Z_{k|k} A^{\mathrm{T}} - Q) G_k^{\mathrm{T}}, \tag{26}$$

$$G_k = Z_{k|k} A^{\mathrm{T}} (A Z_{k|k} A^{\mathrm{T}} + Q)^{-1}, \tag{27}$$

and $z_{k|k}$ and $Z_{k|k}$ are the mean and covariance matrix of the filtering posterior $p(z_k | y_{1:k})$.

Proof: The proof is mostly similar to the proof of [29, Theorem 8.2]. First, assume that

$$z_{k+1} | y_{1:K} \sim \mathcal{N}(z_{k+1|K}, Z_{k+1|K}). \tag{28}$$

for some $k < K$. The joint conditional distribution of z_k and z_{k+1} is then

$$p(z_k, z_{k+1} | y_{1:k}) = p(z_{k+1} | z_k)\, p(z_k | y_{1:k}) \quad || \text{ Markovianity assumption} \tag{29}$$

$$= \mathcal{N}(z_{k+1}; A z_k, Q)\, \mathcal{N}(z_k; z_{k|k}, Z_{k|k}) \tag{30}$$

$$= \mathcal{N}\left(\begin{bmatrix} z_k \\ z_{k+1} \end{bmatrix}; \begin{bmatrix} z_{k|k} \\ A z_{k|k} \end{bmatrix}, \begin{bmatrix} Z_{k|k} & Z_{k|k} A^{\mathrm{T}} \\ A Z_{k|k} & A Z_{k|k} A^{\mathrm{T}} + Q \end{bmatrix} \right), \tag{31}$$

so by the conditioning rule of the multivariate normal distribution

$$p(z_k | z_{k+1}, y_{1:k}) = \mathcal{N}(z_k; z_{k|k} + G_k(z_{k+1} - A z_{k|k}), Z_{k|k} - Z_{k|k} A^{\mathrm{T}} (A Z_{k|k} A^{\mathrm{T}} + Q)^{-1} A Z_{k|k}) \tag{32}$$

$$= \mathcal{N}(z_k; z_{k|k} + G_k(z_{k+1} - A z_{k|k}), Z_{k|k} - G_k(A Z_{k|k} A^{\mathrm{T}} + Q) G_k^{\mathrm{T}}). \tag{33}$$

We use this formula in

$$
\begin{aligned}
p(z_k, z_{k+1}|y_{1:K}) &= p(z_k|z_{k+1}, y_{1:K})\, p(z_{k+1}|y_{1:K}) && (34)\\
&= p(z_k|z_{k+1}, y_{1:k})\, p(z_{k+1}|y_{1:K}) \quad || \text{Markovianity assumption} && (35)\\
&= \mathcal{N}(z_k; z_{k|k} + G_k(z_{k+1} - Az_{k|k}), Z_{k|k} - G_k(AZ_{k|k}A^{\mathrm{T}} + Q)G_k^{\mathrm{T}})\, \mathcal{N}(z_{k+1}|z_{k+1|K}, Z_{k+1|K}) && (36)\\
&= \mathcal{N}\left(\begin{bmatrix} z_k \\ z_{k+1} \end{bmatrix}; \begin{bmatrix} z_{k|k} + G_k(z_{k+1|K} - Az_{k|k}) \\ \bullet \end{bmatrix}, \begin{bmatrix} G_k Z_{k+1|K} G_k^{\mathrm{T}} + Z_{k|k} - G_k(AZ_{k|k}A^{\mathrm{T}} + Q)G_k^{\mathrm{T}} & \bullet \\ \bullet & \bullet \end{bmatrix} \right) && (37)\\
&= \mathcal{N}\left(\begin{bmatrix} z_k \\ z_{k+1} \end{bmatrix}; \begin{bmatrix} z_{k|k} + G_k(z_{k+1|K} - Az_{k|k}) \\ \bullet \end{bmatrix}, \begin{bmatrix} Z_{k|k} + G_k(Z_{k+1|K} - AZ_{k|k}A^{\mathrm{T}} - Q)G_k^{\mathrm{T}} & \bullet \\ \bullet & \bullet \end{bmatrix} \right), && (38)
\end{aligned}
$$

so

$$
\begin{aligned}
p(z_k|y_{1:K}) &= \mathcal{N}(z_k; z_{k|k} + G_k(z_{k+1|K} - Az_{k|k}), Z_{k|k} + G_k(Z_{k+1|K} - AZ_{k|k}A^{\mathrm{T}} - Q)G_k^{\mathrm{T}}) && (39)\\
&= \mathcal{N}(z_k; z_{k|K}, Z_{k|K}). && (40)
\end{aligned}
$$

Because $z_K|y_{1:K} \sim \mathcal{N}(z_{K|K}, Z_{K|K})$, and because (28) implies (40), the statement holds by the induction argument. ■

B. Derivations for q_Λ

Eq. (12b) gives

$$
\log q_\Lambda(\Lambda_{1:K}) = \sum_{k=1}^{K} \left\{ \mathop{\mathbb{E}}_{q_{xu}} \left[\log \mathcal{N}(y_k; Cx_k + \Delta u_k, \Lambda_k^{-1} R) + \log \mathcal{N}_+(u_k; 0, \Lambda_k^{-1}) \right] \right\} + \sum_{k=1}^{K} \sum_{i=1}^{n_y} \log \mathcal{G}\left([\Lambda_k]_{ii}; \frac{\nu_i}{2}, \frac{\nu_i}{2} \right) + c. \tag{41}
$$

Therefore, $q_\Lambda(\Lambda_{1:K}) = \prod_{k=1}^{K} q_\Lambda(\Lambda_k)$ where

$$
\begin{aligned}
\log q_\Lambda(\Lambda_k) = &-\frac{1}{2} \mathop{\mathbb{E}}_{q_{xu}} \left[\mathrm{tr}\{(y_k - Cx_k - \Delta u_k)(y_k - Cx_k - \Delta u_k)^{\mathrm{T}} R^{-1} \Lambda_k\} \right] \\
&- \frac{1}{2} \mathop{\mathbb{E}}_{q_{xu}} \left[\mathrm{tr}\{u_k u_k^{\mathrm{T}} \Lambda_k\} \right] + \sum_{i=1}^{n_y} \left(\frac{\nu_i}{2} \log[\Lambda_k]_{ii} - \frac{\nu_i}{2} [\Lambda_k]_{ii} \right) + c && (42)\\
= &-\frac{1}{2} \mathrm{tr} \left\{ \left((y_k - Cx_{k|K} - \Delta u_{k|K})(y_k - Cx_{k|K} - \Delta u_{k|K})^{\mathrm{T}} + \begin{bmatrix} C & \Delta \end{bmatrix} Z_{k|K} \begin{bmatrix} C^{\mathrm{T}} \\ \Delta^{\mathrm{T}} \end{bmatrix} \right) R^{-1} \Lambda \right\} && (43)\\
&- \frac{1}{2} \mathrm{tr} \left\{ (u_{k|K} u_{k|K}^{\mathrm{T}} + U_{k|K}) \Lambda_k \right\} + \sum_{i=1}^{n_y} \left(\frac{\nu_i}{2} \log[\Lambda_k]_{ii} - \frac{\nu_i}{2} [\Lambda_k]_{ii} \right) + c && (44)\\
= &\sum_{i=1}^{n_y} \left(\frac{\nu_i}{2} \log[\Lambda_k]_{ii} - \frac{\nu_i + [\Psi_k]_{ii}}{2} [\Lambda_k]_{ii} \right) + c && (45)
\end{aligned}
$$

where

$$
\Psi_k = (y_k - Cx_{k|K} - \Delta u_{k|K})(y_k - Cx_{k|K} - \Delta u_{k|K})^{\mathrm{T}} R^{-1} + \begin{bmatrix} C & \Delta \end{bmatrix} Z_{k|K} \begin{bmatrix} C^{\mathrm{T}} \\ \Delta^{\mathrm{T}} \end{bmatrix} R^{-1} + u_{k|K} u_{k|K}^{\mathrm{T}} + U_{k|K}. \tag{46}
$$

Therefore,

$$
q_\Lambda(\Lambda_k) = \prod_{i=1}^{n_y} \mathcal{G}\left([\Lambda_k]_{ii}; \frac{\nu_i}{2} + 1, \frac{\nu_i + [\Psi_k]_{ii}}{2} \right). \tag{47}
$$

Note that only the diagonal elements of the matrix Ψ_k are required. In the derivations of Section A-A, $\Lambda_{k|K} \triangleq \mathbb{E}_{q_\Lambda}[\Lambda_k]$ is required. $\mathbb{E}_{q_\Lambda}[\Lambda_k]$ is a diagonal matrix with the diagonal elements

$$
[\Lambda_{k|K}]_{ii} = \frac{\nu_i + 2}{\nu_i + [\Psi_k]_{ii}}. \tag{48}
$$

APPENDIX B
DERIVATIONS FOR THE FILTER

Suppose that at time index k the measurement vector y_k is available, and the prediction PDF $p(x_k|y_{1:k-1})$ is

$$p(x_k|y_{1:k-1}) = \mathcal{N}(x_k; x_{k|k-1}, P_{k|k-1}). \tag{49}$$

Then, using Bayes' theorem the joint filtering posterior PDF is

$$\begin{aligned} p(x_k, u_k, \Lambda_k|y_{1:k}) &\propto p(y_k, x_k, u_k, \Lambda_k|y_{1:k-1}) && (50)\\ &= p(y_k|x_k, u_k, \Lambda_k)\, p(x_k|y_{1:k-1})\, p(u_k|\Lambda_k)\, p(\Lambda_k) && (51)\\ &= \mathcal{N}(y_k; Cx_k + \Delta u_k, \Lambda_k^{-1}R)\, \mathcal{N}(x_k; x_{k|k-1}, P_{k|k-1})\, \mathcal{N}_+(u_k; 0, \Lambda_k^{-1}) \prod_{i=1}^{n_y} \mathcal{G}\left([\Lambda_k]_{ii}; \frac{\nu_i}{2}, \frac{\nu_i}{2}\right). && (52)\end{aligned}$$

This posterior is not analytically tractable. We seek an approximation in the form

$$p(x_k,u_k, \Lambda_k|y_{1:k}) \approx q_{xu}(x_k, u_k)\, q_\Lambda(\Lambda_k). \tag{53}$$

The VB solutions $\hat{q}_{xu}$ and $\hat{q}_\Lambda$ can be obtained by cyclic iteration of

$$\log q_{xu}(x_k, u_k) \leftarrow \mathop{\mathbb{E}}_{q_\Lambda}[\log p(y_k, x_k, u_k, \Lambda_k|y_{1:k-1})] + c_{xu} \tag{54a}$$

$$\log q_\Lambda(\Lambda_k) \leftarrow \mathop{\mathbb{E}}_{q_{xu}}[\log p(y_k, x_k, u_k, \Lambda_k|y_{1:k-1})] + c_\Lambda \tag{54b}$$

where the expected values on the right hand sides of (54) are taken with respect to the current q_{xu} and q_Λ, and c_{xu} and c_Λ are constants with respect to the variables $\left[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}\right]$ and Λ_k, respectively [10, Chapter 10] [17]. In sections B-A and B-B the derivations for the variational solution (54) are given. For brevity all constant values are denoted by c in the derivations. The logarithm of the joint filtering posterior which is needed for the derivations is given by

$$\begin{aligned} \log p(y_k, x_k, u_k, \Lambda_k|y_{1:k-1}) = &-\frac{1}{2}(y_k - Cx_k - \Delta u_k)^{\mathrm{T}} R^{-1} \Lambda_k (y_k - Cx_k - \Delta u_k)\\ &-\frac{1}{2}(x_k - x_{k|k-1})^{\mathrm{T}} P_{k|k-1}^{-1}(x_k - x_{k|k-1})\\ &-\frac{1}{2}u_k^{\mathrm{T}}\Lambda_k u_k + \sum_{i=1}^{n_y}\left(\frac{\nu_i}{2}\log[\Lambda_k]_{ii} - \frac{\nu_i}{2}[\Lambda_k]_{ii}\right) + c,\ u_k \geq 0, \end{aligned} \tag{55}$$

where $u_k \geq 0$ means that every component of u_k is non-negative.

A. Derivations for q_{xu}

Using equation (54a) we obtain

$$\begin{aligned} \log q_{xu}(x_k, u_k) = &-\frac{1}{2}\mathop{\mathbb{E}}_{q_\Lambda}[(y_k - Cx_k - \Delta u_k)^{\mathrm{T}} R^{-1} \Lambda_k (y_k - Cx_k - \Delta u_k)]\\ &-\frac{1}{2}(x_k - x_{k|k-1})^{\mathrm{T}} P_{k|k-1}^{-1}(x_k - x_{k|k-1}) - \frac{1}{2}\mathop{\mathbb{E}}_{q_\Lambda}[u_k^{\mathrm{T}}\Lambda_k u_k] + c && (56)\\ = &-\frac{1}{2}\left(y_k - [C \quad \Delta]\begin{bmatrix} x_k \\ u_k \end{bmatrix}\right)^{\mathrm{T}} R^{-1}\Lambda_{k|k}\left(y_k - [C \quad \Delta]\begin{bmatrix} x_k \\ u_k \end{bmatrix}\right) && (57)\\ &-\frac{1}{2}\left(\begin{bmatrix} x_k \\ u_k \end{bmatrix} - \begin{bmatrix} x_{k|k-1} \\ 0 \end{bmatrix}\right)^{\mathrm{T}} \begin{bmatrix} P_{k|k-1} & 0 \\ 0 & \Lambda_{k|k}^{-1} \end{bmatrix}^{-1} \left(\begin{bmatrix} x_k \\ u_k \end{bmatrix} - \begin{bmatrix} x_{k|k-1} \\ 0 \end{bmatrix}\right),\ u_k \geq 0, && (58)\end{aligned}$$

where $\Lambda_{k|k} \triangleq \mathbb{E}_{q_\Lambda}[\Lambda_k]$ is derived in section B-B. Hence,

$$q_{xu}(x_k, u_k) \propto \mathcal{N}\left(y_k; [C \quad \Delta]\begin{bmatrix} x_k \\ u_k \end{bmatrix}, \Lambda_{k|k}^{-1}R\right)\mathcal{N}\left(\begin{bmatrix} x_k \\ u_k \end{bmatrix}; \begin{bmatrix} x_{k|k-1} \\ 0 \end{bmatrix}, \begin{bmatrix} P_{k|k-1} & 0 \\ 0 & \Lambda_{k|k}^{-1} \end{bmatrix}\right) \cdot [u_k \geq 0], \tag{59}$$

where $[u_k \geq 0]$ is the Iverson bracket. By Kalman filter's [28] measurement update, this becomes

$$q_{xu}(x_k, u_k) = \frac{1}{C}\mathcal{N}\left(\begin{bmatrix} x_k \\ u_k \end{bmatrix}; z'_{k|k}, Z'_{k|k}\right) \cdot [u_k \geq 0], \tag{60}$$

where

$$z'_{k|k} = \begin{bmatrix} x_{k|k-1} \\ 0 \end{bmatrix} + K_k(y_k - Cx_{k|k-1}), \tag{61}$$

$$Z'_{k|k} = (I - K_k \begin{bmatrix} C & \Delta \end{bmatrix}) \begin{bmatrix} P_{k|k-1} & 0 \\ 0 & \Lambda^{-1}_{k|k} \end{bmatrix}, \tag{62}$$

$$K_k = \begin{bmatrix} P_{k|k-1}C^{\mathrm{T}} \\ \Lambda^{-1}_{k|k}\Delta^{\mathrm{T}} \end{bmatrix} (CP_{k|k-1}C^{\mathrm{T}} + \Delta\Lambda^{-1}_{k|k}\Delta^{\mathrm{T}} + \Lambda^{-1}_{k|k}R)^{-1}. \tag{63}$$

The first and second moments $x_{k|k} \triangleq \mathbb{E}_{q_{xu}}[x_k]$, $u_{k|k} \triangleq \mathbb{E}_{q_{xu}}[u_k]$, $Z_{k|k} \triangleq \mathrm{Var}_{q_{xu}}[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}]$, and $U_{k|k} \triangleq \mathrm{Var}_{q_{xu}}[u_k]$ are required in the derivation of q_Λ in Section B-B, and they can be approximated using the recursive truncation algorithm. For the linear–Gaussian time update to be analytically tractable, the marginal distribution $q_{xu}(x_k)$ is approximated as a normal distribution

$$q_{xu}(x_k) = \int q_{xu}(x_k, u_k)\,\mathrm{d}u_k \approx \mathcal{N}(x_{k|k}, P_{k|k}), \tag{64}$$

where $P_{k|k} \triangleq \mathrm{Var}_{q_{xu}}[x_k]$.

B. Derivations for q_Λ

Using equation (54b) we obtain

$$\begin{aligned}
\log q_\Lambda(\Lambda_k) = &-\frac{1}{2}\mathop{\mathbb{E}}_{q_{xu}}[\mathrm{tr}\{(y_k - Cx_k - \Delta u_k)(y_k - Cx_k - \Delta u_k)^{\mathrm{T}} R^{-1}\Lambda_k\}] \\
&-\frac{1}{2}\mathop{\mathbb{E}}_{q_{xu}}[\mathrm{tr}\{u_k u_k^{\mathrm{T}}\Lambda_k\}] + \sum_{i=1}^{n_y}\left(\frac{\nu_i}{2}\log[\Lambda_k]_{ii} - \frac{\nu_i}{2}[\Lambda_k]_{ii}\right) + c && (65)\\
= &-\frac{1}{2}\,\mathrm{tr}\left\{\left((y_k - Cx_{k|k} - \Delta u_{k|k})(y_k - Cx_{k|k} - \Delta u_{k|k})^{\mathrm{T}} + \begin{bmatrix} C & \Delta \end{bmatrix} Z_{k|k} \begin{bmatrix} C^{\mathrm{T}} \\ \Delta^{\mathrm{T}} \end{bmatrix}\right) R^{-1}\Lambda_k\right\} && (66)\\
&-\frac{1}{2}\,\mathrm{tr}\left\{(u_{k|k}u_{k|k}^{\mathrm{T}} + U_{k|k})\Lambda_k\right\} + \sum_{i=1}^{n_y}\left(\frac{\nu_i}{2}\log[\Lambda_k]_{ii} - \frac{\nu_i}{2}[\Lambda_k]_{ii}\right) + c && (67)\\
= &\sum_{i=1}^{n_y}\left(\frac{\nu_i}{2}\log[\Lambda_k]_{ii} - \frac{\nu_i + [\Psi_k]_{ii}}{2}[\Lambda_k]_{ii}\right) + c && (68)
\end{aligned}$$

where

$$\Psi_k = (y_k - Cx_{k|k} - \Delta u_{k|k})(y_k - Cx_{k|k} - \Delta u_{k|k})^{\mathrm{T}} R^{-1} + \begin{bmatrix} C & \Delta \end{bmatrix} Z_{k|k} \begin{bmatrix} C^{\mathrm{T}} \\ \Delta^{\mathrm{T}} \end{bmatrix} R^{-1} + u_{k|k}u_{k[k}^{\mathrm{T}} + U_{k|k} \tag{69}$$

and the moments $x_{k|k} \triangleq \mathbb{E}_{q_{xu}}[x_k]$, $u_{k|k} \triangleq \mathbb{E}_{q_{xu}}[u_k]$, $Z_{k|k} \triangleq \mathrm{Var}_{q_{xu}}[\begin{smallmatrix} x_k \\ u_k \end{smallmatrix}]$, and $U_{k|k} \triangleq \mathrm{Var}_{q_{xu}}[u_k]$ are derived in Section B-A of this report. Therefore,

$$q_\Lambda(\Lambda_k) = \prod_{i=1}^{n_y} \mathcal{G}\left([\Lambda_k]_{ii}; \frac{\nu_i}{2} + 1, \frac{\nu_i + [\Psi_k]_{ii}}{2}\right). \tag{70}$$

Note that only the diagonal elements of the matrix Ψ_k are required. In the derivations of Section B-A $\Lambda_{k|k} \triangleq \mathbb{E}_{q_\Lambda}[\Lambda_k]$ is required. $\mathbb{E}_{q_\Lambda}[\Lambda_k]$ is a diagonal matrix with the diagonal elements

$$[\Lambda_{k|k}]_{ii} = \frac{\nu_i + 2}{\nu_i + [\Psi_k]_{ii}}. \tag{71}$$

REFERENCES

[1] F. Gustafsson and F. Gunnarsson, "Mobile positioning using wireless networks: possibilities and fundamental limitations based on available wireless network measurements," *IEEE Signal Processing Magazine*, vol. 22, no. 4, pp. 41–53, July 2005.

[2] B.-S. Chen, C.-Y. Yang, F.-K. Liao, and J.-F. Liao, "Mobile location estimator in a rough wireless environment using Extended Kalman-based IMM and data fusion," *IEEE Transactions on Vehicular Technology*, vol. 58, no. 3, pp. 1157–1169, March 2009.

[3] M. Kok, J. D. Hol, and T. B. Schön, "Indoor positioning using ultra-wideband and inertial measurements," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 4, 2015.

[4] K. Kaemarungsi and P. Krishnamurthy, "Analysis of WLAN's received signal strength indication for indoor location fingerprinting," *Pervasive and Mobile Computing*, vol. 8, no. 2, pp. 292–316, 2012, special Issue: Wide-Scale Vehicular Sensor Networks and Mobile Sensing.

[5] M. D. Branco and D. K. Dey, "A general class of multivariate skew-elliptical distributions," *Journal of Multivariate Analysis*, vol. 79, no. 1, pp. 99–113, October 2001.

[6] A. Azzalini and A. Capitanio, "Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, vol. 65, no. 2, pp. 367–389, 2003.

[7] A. K. Gupta, "Multivariate skew t-distribution," *Statistics*, vol. 37, no. 4, pp. 359–363, 2003.

[8] H. Nurminen, T. Ardeshiri, R. Piche, and F. Gustafsson, "Robust inference for state-space models with skewed measurement noise," *IEEE Signal Processing Letters*, vol. 22, no. 11, pp. 1898–1902, Nov 2015.

[9] H. Nurminen, T. Ardeshiri, R. Piché, and F. Gustafsson, "A NLOS-robust TOA positioning filter based on a skew-t measurement noise model," in *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, October 2015, pp. 1–7.

[10] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2007.

[11] M. P. Wand, J. T. Ormerod, S. A. Padoan, and R. Frühwirth, "Mean field variational Bayes for elaborate distributions," *Bayesian Analysis*, vol. 6, no. 4, pp. 847–900, 2011.

[12] S. K. Sahu, D. K. Dey, and M. D. Branco, "Erratum: A new class of multivariate skew distributions with applications to Bayesian regression models," *Canadian Journal of Statistics*, vol. 37, no. 2, pp. 301–302, 2009.

[13] ——, "A new class of multivariate skew distributions with applications to Bayesian regression models," *Canadian Journal of Statistics*, vol. 31, no. 2, pp. 129–150, 2003.

[14] T.-I. Lin, "Robust mixture modeling using multivariate skew t distributions," *Statistics and Computing*, vol. 20, pp. 343–356, 2010.

[15] S. X. Lee and G. J. McLachlan, "EMMIXuskew: An R package for fitting mixtures of multivariate skew t distributions via the EM algorithm," *Journal of Statistical Software*, vol. 55, no. 12, pp. 1–22, November 2013.

[16] T. M. Cover and J. Thomas, *Elements of Information Theory*. John Wiley and Sons, 2006.

[17] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, "The variational approximation for Bayesian inference," *IEEE Signal Processing Magazine*, vol. 25, no. 6, pp. 131–146, Nov. 2008.

[18] M. J. Beal, "Variational algorithms for approximate Bayesian inference," Ph.D. dissertation, Gatsby Computational Neuroscience Unit, University College London, 2003.

[19] G. Tallis, "The moment generating function of the truncated multi-normal distribution," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 23, no. 1, pp. 223–119, 1961.

[20] A. Genz, "Numerical computation of rectangular bivariate and trivariate normal numerical computation of rectangular bivariate and trivariate normal and t probabilities," *Statistics and Computing*, vol. 14, pp. 251–260, 2004.

[21] A. Genz and F. Bretz, "Comparison of methods for the computation of multivariate t probabilities," *Journal of Computational and Graphical Statistics*, vol. 11, no. 4, pp. 950–971, 2002.

[22] T. Perälä and S. Ali-Löytty, "Kalman-type positioning filters with floor plan information," in *6th International Conference on Advances in Mobile Computing and Multimedia (MoMM)*. New York, NY, USA: ACM, 2008, pp. 350–355.

[23] D. J. Simon and D. L. Simon, "Constrained Kalman filtering via density function truncation for turbofan engine health estimation," *International Journal of Systems Science*, vol. 41, no. 2, pp. 159–171, 2010.

[24] H. E. Rauch, C. T. Striebel, and F. Tung, "Maximum Likelihood Estimates of Linear Dynamic Systems," *Journal of the American Institute of Aeronautics and Astronautics*, vol. 3, no. 8, pp. 1445–1450, August 1965.

[25] J. M. Dow, R. Neilan, and C. Rizos, "The international GNSS service in a changing landscape of global navigation satellite systems," *Journal of Geodesy*, vol. 83, no. 7, p. 689, February 2009.

[26] Y. Bar-Shalom, R. X. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation, Theory Algorithms and Software*. John Wiley & Sons, 2001.

[27] R. Piché, S. Särkkä, and J. Hartikainen, "Recursive outlier-robust filtering and smoothing for nonlinear systems using the multivariate Student-t distribution," in *IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, September 2012.

[28] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME–Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[29] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge, UK: Cambridge University Press, 2013.